

**Abstract of the Disclosure****NUMA SYSTEM RESOURCE DESCRIPTORS INCLUDING  
PERFORMANCE CHARACTERISTICS**

5 A method and computer system for efficiently accessing resources in a  
 multiprocessor computer system. Multiple resources are grouped into nodes and system  
 resource descriptors are maintained as data structures. The method incorporates  
 traversing a data structure to efficiently allocate resources within a grouping of nodes.  
 Each node in the system is assigned a node identifying number. This number identifies a  
 node location within a multiprocessor and is used to determine latency between nodes,  
 10 either through an average latency table or a system interconnect connection table. The  
 data structure comprises secondary data structures therein for storing processor, bus,  
 memory and shared cache information. The data structure includes pointers to each of the  
 secondary data structures. In addition, each node or grouping of nodes may include  
 subnodes. As such, the system provides for a method of recursively accessing additional  
 15 data structure levels for each level of nodes and/or subnodes in the system. In addition to  
 the data structure and the secondary data structure, the system includes a memory map in  
 the form of a data structure. This data structure stores information as to which node a  
 given block of memory address space is located in. In addition, there is a flag for each  
 memory block indicating whether the memory address is parity or error correcting code  
 20 (ECC), or neither parity nor ECC protected. This flag is in the form of two bits for  
 indicating the level of memory protection provided to a specific block of memory address.  
 Accordingly, storing the system architecture information in a data structure format  
 provides for an efficient method of accessing system resources and determining latency for  
 specific processes and access patterns.